

Social Dynamics of Science

Xiaoling Sun,^{1,2} Jasleen Kaur,² Staša Milojević,³
Alessandro Flammini,² Filippo Menczer²

¹ Department of Computer Science and Technology
Dalian University of Technology, China

² Center for Complex Networks and Systems Research
School of Informatics and Computing

³ School of Library and Information Science
Indiana University, Bloomington, USA

Abstract

The birth and decline of disciplines are critical to science and society. However, no quantitative model to date allows us to validate competing theories of whether the emergence of scientific disciplines drives or follows the formation of social communities of scholars. Here we propose an agent-based model based on a *social dynamics of science*, in which the evolution of disciplines is guided mainly by the social interactions among scientists. We find that such a social theory can account for a number of stylized facts about the relationships between disciplines, authors, and publications. These results provide strong quantitative support for the key role of social interactions in shaping the dynamics of science. A “science of science” must gauge the role of exogenous events, such as scientific discoveries and technological advances, against this purely social baseline.

1 Introduction

Understanding the dynamics of science as a human endeavor — the birth, evolution, and decline of disciplines — is of critical importance for allocating resources and planning toward positive societal impact. For example, the emergence of new fields such as bioinformatics, nanophysics, information technology, quantum computing, and data science promises “converging technologies” with unparalleled potential to influence our lives. This paper is about modeling the dynamic evolution of scientific disciplines.

Efforts to describe, explain and predict different aspects of science have intensified in recent years [5, 3, 24] spanning a wide range of theoretical, mathematical, statistical and computational approaches. A number of models of the dynamics of science have explored the continually changing disciplinary relations. Whether the new disciplines or fields are the results of branching of the old ones

due to growth and new discoveries [23, 15], or go through a “specialization-fragmentation-hybridization” cycle [7], or grow through the synthesis of elements of pre-existing ones [8], these models point to the self-organizing development of science exhibiting growth and emergent behavior [20, 25, 26].

Kuhn’s *cognitive* theory emphasizes the role of observations not explained by previous paradigms [12]. Other scholars [6, 27] emphasize the formation of *social groups* of scientists as the driving force behind the evolution of disciplines. These models, however, are difficult to validate empirically or lack explanations of the processes leading to the empirical patterns they describe.

Agent-based models [22] allow one to generate macroscopic predictions from micro-level mechanisms guiding the behavior of individuals, thus providing a powerful approach to model the emergence of disciplines. While this approach has been used to create models of science dynamics [9, 4, 28], the focus was primarily on coauthorship, publication, and citation behavior rather than the emergence of disciplines.

Quantitative work on modeling the emergence of disciplines is lacking to date, owing in part to the difficulty of formally defining the notion of scientific field, and the consequent sparsity of data to inform and validate models. Here we offer a first quantitative baseline model to explore the consequences of assuming a purely social basis of science dynamics, without explicit references to exogenous events such as scientific discoveries. In our model agents represent scholars who choose their collaborators, while groups of collaborating scholars represent scientific disciplines [21]. The key idea behind our model is that new scientific fields emerge from splitting and merging of these social communities. Our model thus defines a *social dynamics of science*, in which the birth and evolution of disciplines is guided mainly by the social interactions among scientists. We find that such a social theory can account for several stylized facts about the relationships between disciplines, authors, and publications.

2 Model Description

The critical assumption of our model is the correspondence between the social dynamics of scholar communities and the evolution of scientific disciplines. To illustrate this intuition, let us look at the coauthorship network for papers published by the American Physical Society (APS). Using journals as proxies for scholarly communities, we can track the changes in community structure over time. Fig. 1 plots the modularity [19] of the partition induced by the journals; higher values indicate a more clustered structure (see Methods).

We observe noticeable changes in modularity around the introduction of new journals. Some of these changes suggest a scenario in which a new field emerges (e.g., quantum mechanics in the late 1920’s), and a new journal captures the corresponding scholar community, leading to an increase in modularity. Interdisciplinary interactions across established areas lead to a decrease in modularity (e.g., prior to the introduction of *Physical Review E* in the 1990’s). These observations motivate a community detection approach to model the evolution of

disciplines.

In the proposed model, which we call SDS (for Social Dynamics of Science), we build a social network of collaborations whose nodes are authors, linked by coauthored papers as illustrated in Fig. 2(a). Each author is represented by a list of disciplines indicating the scientific fields they have been working on, and every discipline has a list of papers. Similarly, each link is represented by a list of disciplines with associated papers describing the collaborations between two authors.

There are three elements in the SDS model: papers, authors, and disciplines. The social network starts with one author writing one paper in one discipline. The network then evolves as new authors join, new papers are written, and new disciplines emerge over time.

At every time step, a new paper is added to the network. Its first author is chosen uniformly at random, so every author has a chance to publish a paper. In modeling the choice of collaborators, we aim to capture a few basic intuitions: (i) authors who have collaborated before are likely to do so again; (ii) authors with common collaborators are likely to collaborate with each other; (iii) it is easier to choose collaborators with similar than dissimilar background; and (iv) authors with many collaborations have higher probability to gain additional ones [14, 1]. We model these behaviors through a biased random walk [13], illustrated in Fig. 2(b). The length of the random walk determines the number of coauthors. At each step in the walk, the author visits a node i (starting with itself) and decides to stop with probability p_w , or to search for additional authors with probability $1 - p_w$. In the latter case a neighbor j is selected as coauthor according to the transition probability:

$$P_{ij} = \frac{w_{ij}}{\sum_k w_{ik}} \quad (1)$$

where w_{ij} is the weight of the edge connecting authors i and j , that is, the number of papers that i and j have written together. Note that the walk may result in a single author.

We propose a simple mechanism to model knowledge diffusion through collaboration in SDS: when authors write a paper together, they all contribute their knowledge. Therefore, a paper inherits the union of the author disciplines as topics. However, the discipline that is shared by the majority of authors is selected as the main topic of the paper (say, the publication venue) and diffuses across all the authors. Through the collaboration, authors acquire knowledge of and membership in this area.

At every time step, with probability p_n , we add a new author to the network with the new paper. The parameter p_n regulates the ratio of papers to authors. The new author is the first author of the new paper. To generate other collaborators, an existing author is first selected uniformly at random as the first coauthor. Then the random walk procedure is followed to pick additional collaborators. The new author acquires the main topic of the paper.

We introduce a novel mechanism to model the evolution of disciplines by splitting and merging communities in the social collaboration network. The

idea, motivated by the earlier observations from the APS data, is that the birth or decline of a discipline should correspond to an increase in the modularity of the network. Two such events may occur at each time step with probability p_d . The process is illustrated in Fig. 3.

For a *split* event we select a random discipline with its coauthor network and decide whether a new discipline should emerge from a subset of this community. We partition the coauthor network into two clusters (see Methods). If the modularity of the partition is higher than that of the single discipline, there are more collaborations within each cluster than across the two. We then split the smaller community as a new discipline. In this case the papers whose authors are all in the new community are relabeled to reflect the emergent discipline. Borderline papers with authors in both old and new disciplines are labeled according to the discipline of the majority of authors. Some authors may as a result belong to both old and new discipline.

For a *merge* event we randomly select two disciplines with at least one common author. If the modularity obtained by merging the two groups is higher than that of the partitioned groups, the collaborations across the two communities are stronger than those within each one. The two are then merged into a single new discipline. In this case all the papers in the two old disciplines are relabeled to reflect the new one.

3 Results

To evaluate the predictive power of the SDS model we consider a number of *stylized facts*, i.e., broad empirical observations that describe essential characteristics of the dynamic relationships between disciplines, scholars, and publications. Our model provides an explanation for the evolution of scientific fields if it can reproduce these empirical observations. The complex interactions of a changing group of scientists, their artifacts, and their disciplinary aggregations can be captured by the broad empirical distributions of six quantitative descriptors: the number of authors per paper A_P (collaboration size); the number of papers per author P_A (scholar productivity); the number of authors per discipline A_D (discipline popularity); the number of disciplines per author D_A (scholar interdisciplinary effort); the number of papers per discipline P_D (discipline productivity); and the number of disciplines per paper D_P (publication breadth).

To validate the SDS model, one would ideally require a real-world dataset mapping the three-way relationships between scholars, publications, and disciplines. Unfortunately, to the best of our knowledge, no such dataset is publicly available. As an alternative, we adopt three large datasets that each map a binary projection of these relationships: *NanoBank* [30] to validate the relationship between authors and papers, *Scholarometer* [11] to study the relationship between authors and disciplines, and *Bibsonomy* [2] to analyze the relationship between papers and disciplinary topics. The datasets are described in the Methods section. The parameters p_n , p_w , and p_d of our model are tuned to fit the

quantitative descriptors of each dataset (see Methods).

Fig. 4 presents a close match between the real data and the predictions of our model. SDS reproduces the stylized facts about the relationships between scholars, publications, and disciplines, characterized by these six distributions. The exponential distribution of A_P is captured by the random walk process. The broad distribution of scholar productivity P_A is well accounted for by the bias in the random walk, which incorporates a kind of preferential attachment mechanism regulated by prior collaborations. The distributions of discipline popularity A_D and productivity P_D also display heavy tails, which cannot be attributed to a specific mechanism in the model; they emerge from the non-trivial interactions between (i) merging and splitting of the discipline communities and (ii) knowledge diffusion from the collaborations. The prediction is not as good for D_A : our model produces a relatively large number of highly interdisciplinary authors. One could correct this effect, for example, by requiring more than one paper in a discipline as a condition for membership. However, this would require an additional parameter and thus a more complicated model. Finally, The distribution of publication D_P shows that there is a continuum in the breadth of papers, rather than a sharp separation between disciplinary and interdisciplinary work.

These results focus on the relationships between disciplines, authors, and papers, for which there is little prior quantitative analysis. The coauthor network, on the other hand, has been studied extensively in the past [16, 17]. As shown in Fig. 5, the SDS model generates coauthor networks whose long-tailed degree distributions are consistent with the empirical data, as well as with those in the literature.

4 Conclusions

We introduced an agent-based model to simulate the evolution of science as a process driven only by social dynamics. The model captures for the first time major stylized facts about the complex socio-cognitive interactions of a changing group of scholars, publications, and scientific communities. The SDS model is relatively simple when one considers the complexity of the science dynamics process being studied, yet powerful in its capability to reproduce the emergence of patterns similar to those observed in three real datasets about scientific production and fields. This provides us with strong quantitative support for the key role of social dynamics in shaping the birth, evolution, and decline of scientific disciplines. Future “science of science” studies will have to gauge the role of scientific discoveries, technological advances, and other exogenous events in the emergence of new disciplines against this purely social baseline.

5 Methods

Modularity [19] measures the strength of a network partition into clusters of nodes. It compares the number of edges falling within groups with the expected number in an equivalent network from a null model with the same degree sequence but shuffled edges. Larger values indicate stronger community structure. Here we consider the weighted extension of modularity. Let w_{ij} be the weight of an edge (number of coauthored papers) between nodes i and j , and W_{ij} its expected value. The weighted modularity is defined as:

$$Q = \frac{1}{2m} \sum_{ij} [w_{ij} - W_{ij}] \delta(g_i, g_j) \quad (2)$$

where $\delta(g_i, g_j) = 1$ if $g_i = g_j$ (i and j are in the same group) and 0 otherwise; m is the sum of all edge weights in the network. W_{ij} is computed as:

$$W_{ij} = \frac{s_i s_j}{2m} \quad (3)$$

where s_i is the *strength* or weighted degree of node i , $s_i = \sum_j w_{ij}$.

When splitting disciplines, in practice, we use the leading eigenvector method [18] based on the (non-weighted) modularity matrix, as an efficient and effective algorithm to cluster a coauthor network into two groups.

The APS dataset (Fig. 1) was made available by the American Physical Society (publish.aps.org/datasets/). We consider the papers appearing in eight journals during the period of 1913-2000: *Physical Review* (PR) 1913-1955, *Review of Modern Physics* (RMP) 1929-2000, *Physical Review Letters* (PRL) 1958-2000, *Physical Review A, B, C, D* (PRA-D) 1970-2000, and *Physical Review E* (PRE) 1993-2000.

The SDS model is validated against three datasets:

NanoBank (version Beta 1, released on May 2007) [30] is a digital library of bibliographic data on articles, patents and grants in the field of nanotechnology. A set of nanotechnology-related articles in NanoBank has been selected from the Science Citation Index Expanded, Social Sciences Citation Index, and Arts and Humanities Citation Index produced by the Institute for Scientific Information (now Thomson Reuters). Two document selection criteria have been used in the creation of NanoBank [29]: (i) articles that contain some of the 379 terms identified by subject specialists as being “nano-specific,” and (ii) articles based on a probabilistic procedure for the automatic identification of terms. The database covers a 35-year period (1970-2004). This dataset is used to validate the relationship between authors and papers.

Scholarometer (scholarometer.indiana.edu) is a social tool for scholarly services developed at Indiana University, with the goal of exploring the crowdsourcing approach for disciplinary annotations and cross-disciplinary impact metrics [10, 11]. Users provide discipline annotations (tags) for

queried authors, which in turn are used to compare author impact across disciplinary boundaries. The data collected by Scholarometer is available via an open API. We use this data to study the relationship between authors and disciplines.

Bibsonomy (www.bibsonomy.org) is a system for sharing bookmarks and lists of literature [2]. Users annotate papers with tags. The dataset is free for research purposes. We downloaded a dump as of 2012-01-01 to analyze the relationship between papers and disciplines. To filter noise from many junk annotations, we removed the tags associated with fewer than 3 papers or more than 6,000 papers.

The SDS model has three parameters: p_n controls the number of papers per author; p_w controls the number of authors per paper; finally, p_d is the frequency of network split and merge events, and controls the number of disciplines. To generate predictions from numerical simulations of the model, we first tune these parameters to fit the properties of the three empirical datasets individually. Table 1 reports the main properties of the datasets and the matching model parameters. As shown in Table 2, the SDS model is capable of approximating the basic statistics of the empirical data.

Acknowledgments

The work presented in this paper was performed while Xiaoling Sun was visiting the Center for Complex Networks and Systems Research (cnets.indiana.edu) at the Indiana University School of Informatics and Computing. Thanks to Diep Thi Hoang, Mohsen JafariAsbagh, and Lino Possamai for Scholarometer system development and helpful discussions, and John McCurley for editing assistance. We acknowledge support from Hongfei Lin at Dalian University of Technology, the China Scholarship Council, the Lilly Endowment, and NSF (award IIS-0811994) for funding the computing infrastructure that hosts the Scholarometer service.

References

- [1] A. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [2] D. Benz, A. Hotho, R. Jäschke, B. Krause, F. Mitzlaff, C. Schmitz, and G. Stumme. The social bookmark and publication management system bibsonomy. *The VLDB Journal*, 19(6):849–875, 2010.
- [3] K. Börner, W. Glänzel, A. Scharnhorst, and P. V. den Besselaar. Modeling science: studying the structure and dynamics of science. *Scientometrics*, 89(1):347–348, 2011.

- [4] K. Börner, J. Maru, and R. Goldstone. The simultaneous evolution of author and paper networks. *Proc. Nat. Acad. Sci. USA*, 101:5266–5273, 2004.
- [5] K. Börner and A. Scharnhorst. Visual conceptualizations and models of science. *J. Informetrics*, 3(3):161–172, 2009.
- [6] D. Crane. *Invisible colleges: Diffusion of knowledge in scientific communities*. University of Chicago Press, 1972.
- [7] M. Dogan and R. Pahre. *Creative marginality: Innovation at the intersections of social sciences*. Westview Press, 1990.
- [8] H. Etzkowitz and L. Leydesdorff. The dynamics of innovation: from national systems and. *Research Policy*, 29(2):109–123, 2000.
- [9] N. Gilbert. A simulation of the structure of academic science. *Sociological Research Online*, 2, 1997.
- [10] D. T. Hoang, J. Kaur, and F. Menczer. Crowdsourcing scholarly data. In *Proc. Web Science Conference: Extending the Frontiers of Society On-Line (WebSci)*, 2010.
- [11] J. Kaur, D. T. Hoang, X. Sun, L. Possamai, M. JafariAsbagh, S. Patil, and F. Menczer. Scholarometer: A Social Framework for Analyzing Impact across Disciplines. *PLoS ONE*, 7(9):e43235, 09 2012.
- [12] T. S. Kuhn. *The structure of scientific revolutions*. University of Chicago Press, 1996.
- [13] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdos is Eighty*, 2(1):1–46, 1993.
- [14] R. Merton. The Matthew effect in science. *Science*, 159(3810):56–63, 1968.
- [15] M. Mulkay. Three models of scientific development. *The Sociological Review*, 23(3):509–526, 1975.
- [16] M. E. J. Newman. The structure of scientific collaboration networks. *Proc. Nat. Acad. Sci. USA*, 98(2):404–409, 2001.
- [17] M. E. J. Newman. Coauthorship networks and patterns of scientific collaboration. *Proc. Nat. Acad. Sci. USA*, 101:5200–5205, 2004.
- [18] M. E. J. Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, 74(3):036104, Sept. 2006.
- [19] M. E. J. Newman. Modularity and community structure in networks. *Proc. Nat. Acad. Sci. USA*, 103(23):8577–8582, 2006.

- [20] E. C. M. Noyons and A. F. J. van Raan. Monitoring scientific developments from a dynamic perspective: Self-organized structuring to map neural network research. *Journal of the American Society for Information Science*, 49(1):68–81, 1998.
- [21] G. Palla, A. Barabasi, and T. Vicsek. Quantifying social group evolution. *Nature*, 446(7136):664–667, 2007.
- [22] N. Payette. Agent-based models of science. *Models of Science Dynamics*, pages 127–157, 2012.
- [23] D. Price. *Little science, big science*. Columbia University Press, 1986.
- [24] A. Scharnhorst, K. Borner, and P. Besselaar. *Models of Science Dynamics: Encounters Between Complexity Theory and Information Sciences*. Springer Verlag, 2012.
- [25] A. F. J. van Raan. Fractal dimension of co-citations. *Nature*, 347:626, 1990.
- [26] A. F. J. van Raan. On growth, ageing, and fractal differentiation of science. *Scientometrics*, 47(2):347–362, 2000.
- [27] C. Wagner. *The new invisible college: Science for development*. Brookings Institution Press, 2008.
- [28] C. Watts and N. Gilbert. Does cumulative advantage affect collective learning in science? An agent-based simulation. *Scientometrics*, 89(1):437–463, 2011.
- [29] L. Zucker, M. Darby, J. Furner, R. Liu, and H. Ma. Minerva unbound: Knowledge stocks, knowledge flows and new knowledge production. *Research Policy*, 36(6):850–863, 2007.
- [30] L. G. Zucker and M. R. Darby. Nanobank data description, release 1.0 (beta-test). UCLA Center for International Science, Technology, and Cultural Policy and Nanobank, January 2007.

Figures and Tables




Figure 1: Modularity Q of APS journal-induced scholar communities. For each year t , we build a coauthor network based on the papers published in the 5-year time interval between $t - 2$ and $t + 2$. Such a network snapshot consists only of active authors, who published at least one paper in that time window. If an author published papers in more than one journal, we select the first journal in that period. The grey areas correspond to the introduction of major new journals.

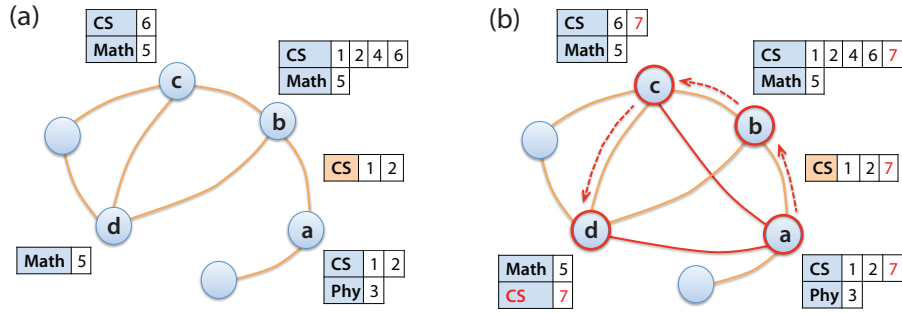


Figure 2: (a) Illustration of the social network structure. Nodes and edges represent authors and their collaborations. They are annotated with lists of (co)authored papers grouped by scientific fields. For example, author b has five papers including four in computer science (CS) and one in Math. Papers 1 and 2 are coauthored with a , papers 5 and 6 with c , and paper 5 with d . Paper 4 is authored by b alone. (b) Illustration of the random walk mechanism to select authors. For the new paper 7, the first author a is chosen randomly and then walks to b and c , stopping at d . These four authors become connected to each other if they have not collaborated before; for example, new edges connect a to c and d . Paper 7 acquires topics CS, Math and Physics (Phy). The main (majority) field of the paper, CS, diffuses across the coauthors, including d who joins this discipline as a result.




Figure 3: Discipline evolution. (a) The coauthor network of discipline D_1 is split into two disciplines D_2 and D_3 . The modularity increases from $Q = 0$ to $Q = 0.4$. The dashed line indicates the partition of the network suggested by the community detection algorithm. Some nodes in the new discipline D_3 have also published papers with authors in D_2 , and therefore belong to both disciplines. (b) Two coauthor networks of disciplines D_4 and D_5 are merged into new discipline D_6 . For authors in both original disciplines, we pick one based on the number of papers published in each discipline. The dashed line shows the resulting partition, with very low modularity $Q = -0.1$. The merged community D_6 has still low, but higher modularity $Q = 0$.

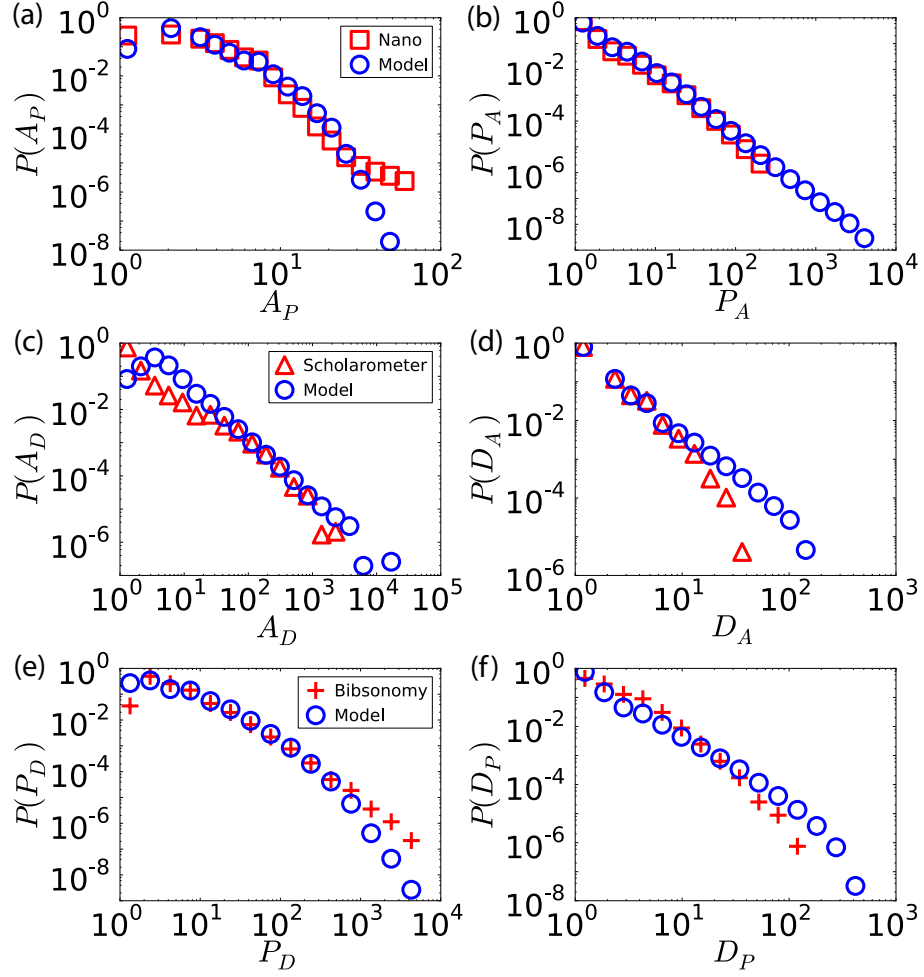


Figure 4: Stylized facts characterizing relationships between authors, papers, and disciplines. We plot the distributions of (a) authors per paper, (b) papers per author, (c) authors per discipline, (d) disciplines per author, (e) papers per discipline, and (f) disciplines per paper. Blue circles represent the SDS predictions, while red symbols represent the empirical data from the three datasets. The results of the model are averaged over 10 runs.




Figure 5: Degree distribution of the coauthor network generated by the SDS model, compared to the empirical distribution from the Bibsonomy dataset. A similar match is also observed for other datasets (not shown). A few papers with more than 100 authors were excluded as they generate an anomaly in the tail; each such paper generates at least 100 nodes with degree at least 100.

Table 1: Dataset properties and tuning of SDS model parameters. For each dataset we run the simulations until the empirical number of papers or authors is reached (shown in bold).

	NanoBank	Scholarometer	Bibsonomy
Number of papers	2.7×10^5	1.4×10^6	2.9×10^5
Number of authors	2.9×10^5	2.2×10^4	3.2×10^5
Number of disciplines	n/a	1.1×10^3	4.4×10^4
p_n	0.90	0.04	0.80
p_w	0.28	0.35	0.71
p_d	n/a	0.01	0.50

Table 2: Basic statistics of empirical datasets compared with SDS model predictions. The reported values are averages, and standard deviations are obtained by 10 realizations of the model.

Quantity	Dataset	Empirical	SDS
A_P	NanoBank	4.006	4.011 ± 0.001
P_A	NanoBank	3.666	4.456 ± 0.005
A_D	Scholarometer	45	60 ± 10
D_A	Scholarometer	2.2	3.5 ± 0.4
P_D	Bibsonomy	24	22 ± 1
D_P	Bibsonomy	3.6	3.3 ± 0.2